Archives available at journals.mriindia.com

Multidisciplinary Journal of Research in Engineering and Technology

ISSN: 2348-6953

Volume 13 Issue 01, 2026

AI Meeting Agent & Automation Platform

[1]Shaikh Mohd Hasan, [2]Sayyad Hamza Fahim, [3]Shrivastav Ujjval Umesh, [4]Rathod Rahul Deepak, [5] Ms. Manasi Churi

[1,2,3,4] Department of Computer Engineering Shree L. R. Tiwari College of Engineering Mira Road, India
[5]Guide, Department of Computer Engineering Shree L. R. Tiwari College of Engineering Mira Road, India
Email: [1] mohd.h.shaikh@slrtce.in, [2]hamza.f.sayyad@slrtce.in, [3]ujjval.u.shrivastav@slrtce.in, [4]rahul.d.rathod@slrtce.in, [5]manasi.churi@slrtce.in

Peer Review Information

Submission: 21 March 2026

Revision: 10 April 2026

Acceptance: 26 April 2026

Keywords

Meeting Summarization, Action Item Detection, Large Language Models, Task Automation, Project Management Tool Integration, Accountability Agent.

Abstract

Meetings are essential, but the work after them is a problem. Someone has to type up notes, pick out action items and enter them into project tools by hand. This manual process leads to mistakes, missed deadlines and tasks that fall through the cracks. Current tools can generate transcripts and summaries, but that's where they stop. They don't create tasks or remind people to follow up. This paper presents an AI Meeting Agent and Automation Platform that fixes these problems. It uses Google's Gemini 1.5 Flash [1] to process meeting transcripts through a five-stage pipeline: secure upload, JSON generation (summary, decisions, action items), user verification and API integration. The platform automatically creates tasks in Trello, Jira and Asana via OAuth 2.0. We also built an Accountability Agent that monitors tasks and sends reminders. Evaluation on 50 transcripts shows 94.2% precision, 91.8% recall and a 79% faster follow-up time. Our work shows that meetings can actually lead to action – less manual work, more accountability.

Introduction

Meetings are central to organizational decision-making, yet the transition from discussion to execution remains a persistent challenge. While teams generate valuable insights and decisions during meetings, converting these into actionable tasks within project management systems is often manual, error-prone, and inefficient. This results in missed responsibilities, unclear ownership, and delayed project timelines.

Recent advances in artificial intelligence, particularly in natural language processing (NLP) and large language models (LLMs), have significantly improved the ability to process and understand human conversations. Transformer-based architectures introduced by Vaswani et al. [16] have revolutionized sequence modeling and serve as the foundation of modern LLMs. Advanced multimodal models such as Gemini [4] demonstrate strong reasoning and contextual understanding capabilities. Additionally, hybrid deep learning approaches combining CNNs, LSTMs, and transformers have proven effective across diverse domains such as healthcare prediction, behavioral modeling, and time-series forecasting [17], [24], [25], [32].

AI techniques have also been successfully applied in highly specialized domains. For instance, adaptive transformer-based frameworks have been used in neurotherapeutic decision support systems [2], while signal processing techniques such as DCT, FFT, and wavelet-based transformations have demonstrated effectiveness in pattern recognition tasks [8], [26]. Similarly, data-driven approaches like Apriori-based optimization models have improved consumer behavior analysis [10]. These cross-domain applications highlight the versatility and scalability of AI models in extracting structured insights from complex, unstructured data.

Existing systems primarily focus on isolated functionalities such as summarization or transcription. Pointer-generator networks [1] improve factual accuracy in summaries, while LLM-based refinement techniques help reduce errors such as irrelevance and hallucination [14]. On-device summarization models like LFM2-2.6B [12] enable privacy-preserving processing without cloud dependency.

Action item detection has also seen progress. Li et al. [3] proposed context-aware models for extracting actionable insights from meeting conversations. Named Entity Recognition (NER) techniques [18] assist in identifying key entities such as people, deadlines, and responsibilities. However, recent evaluation frameworks such as MeetBench-XL [23] highlight that most systems focus on benchmarking and evaluation rather than real-world execution and automation.

Parallel advancements in IoT systems [7], [11], intelligent embedded systems [30], and distributed architectures [15] highlight the feasibility of scalable real-time automation platforms. Security and data integrity are supported through encryption techniques [13] and optimization-driven intelligent systems [21].

However, current solutions lack integration and end-to-end automation. To address this, we propose the **AI Meeting Agent & Automation Platform**, which transforms meeting conversations into actionable tasks and integrates directly with project management tools. The system includes an **Accountability Agent**, inspired by enforcement-based multi-agent frameworks [9] and traceable pipeline architectures [20], ensuring proactive follow-up and task completion.

Literature Review

1. Foundations A. Foundations of Meeting Summarization

Meeting summarization has become essential in modern collaborative environments. Transformer-based architectures [16] enable models to capture contextual relationships effectively, forming the backbone of LLMs such as Gemini [4].

Pointer-generator networks [1] address limitations of abstractive summarization by combining extractive and generative methods, reducing factual inconsistencies. LLM-based refinement techniques further improve summary quality by correcting common errors [14].

Recent advancements include on-device summarization models like LFM2-2.6B [12], which ensure data privacy. Additionally, machine learning optimization techniques [6], [31] contribute to improving summarization robustness and performance.

2. Action Item Detection from Conversations

Extracting actionable tasks from conversations is more complex than summarization. Li et al. [3] introduced context-aware approaches that consider both global and local conversational context.

NER techniques [18] play a crucial role in identifying task-related entities. Broader applications of machine learning in domains such as traffic systems [19] and healthcare analytics [22] demonstrate the effectiveness of AI in extracting meaningful patterns from unstructured data.

3. Frameworks for System Integration and Automation

For real-world deployment, systems must integrate seamlessly with enterprise tools. Cloud-based platforms [5] and distributed systems [15] provide scalable infrastructure for such integrations. IoT-based systems [7], [11] and embedded intelligent devices [29], [30] further enable real-time automation.

Security remains a critical concern, addressed through encryption techniques [13] and intelligent optimization frameworks [21]. Multi-agent systems provide a strong foundation for automation—Tamang and Bora [9] introduced enforcement agents for accountability, while Barrak [20] explored traceability in multi-agent pipelines.

Additional advancements in biometric systems [27], attendance automation [28], and intelligent hardware systems [29] reflect the broader integration of AI into operational workflows.

4. Cross-Domain Machine Learning Applications

Machine learning and artificial intelligence techniques have demonstrated significant success across various domains beyond meeting analysis. In healthcare, transformer-based architectures have been applied for brain activity mapping and intelligent neurotherapeutic decision support [2], while graph-based deep learning models have enabled advanced recommendation systems in medical applications [24]. Similarly, time-series forecasting models using LSTM variants have been effectively used for economic predictions such as GDP growth analysis [25].

In the field of signal and pattern recognition, transform-based feature extraction methods including DCT, FFT, and wavelet transforms have shown strong performance in applications like signature recognition and image processing [8], [26]. These techniques highlight the importance of extracting meaningful representations from raw data—an essential requirement for meeting transcript understanding.
Furthermore, data mining approaches such as Apriori-based optimization models have been widely used in consumer behavior analysis to identify patterns and associations in large datasets [10]. These approaches demonstrate how structured insights can be derived from unstructured or semi-structured data sources.
The success of these cross-domain applications reinforces the adaptability of AI techniques and supports their application in meeting automation systems, where similar challenges of unstructured data understanding, pattern extraction, and decision-making exist.

5. Research Gap Analysis

Despite significant advancements, existing literature reveals key limitations:

- Systems focusing on summarization lack task execution capabilities [1], [4], [12], [14].
- Action item detection methods do not extend to automated task creation [3], [18].
- Multi-agent frameworks are not tailored for meeting-based workflows [9], [20].
- Supporting technologies in IoT, cloud computing, and distributed systems lack domain-specific integration [5], [7], [11], [15], [30].
- Foundational AI and optimization techniques are not applied to complete meeting automation pipelines [6], [16], [31].

Proposed System Methodology

1. System Architecture

Figure 1 shows the five-stage pipeline our AI Meeting Agent follows. The workflow begins with secure transcript upload (Stage 1), followed by AI-driven analysis using Google's Gemini 1.5 Flash model [1] (Stage 2). In Stage 3, once Gemini returns the JSON, we unpack it and shows the user three things: a one-paragraph summary, a list of decisions and a table of action items. Once the user confirms, these action items are automatically sent to Trello, Jira and Asana via OAuth 2.0 APIs (Stage 4). Finally, a backend Accountability Agent runs every 30 minutes to check tasks and remind people through email or Slack messages (Stage 5).

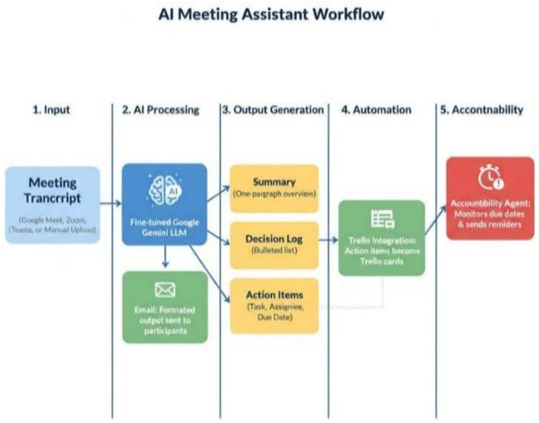

Fig. 1: System Architecture

2. Experimental Setup

For our evaluation, we gathered 50 meeting transcripts from three different domains: 20 came from software development teams, 15 from marketing departments and another 15 from academic research group meetings. Transcripts ranged from 15 to 90 minutes in duration (mean: 42 minutes). Two annotators, working independently, manually reviewed each transcript to identify tasks, assignees and deadlines. Their agreement rate was 92.3%. Precision and recall were calculated as:
We assessed performance using two standard measures: precision (the share of correctly identified items among all system-generated flags) and recall (the share of actual action items that our system captured).

3. Five-Stage Processing Pipeline of the Proposed System

Our system uses a sequential five-stage pipeline:

Stage 1: Input

Users securely upload meeting transcripts in .txt format via a web interface. The proposed system checks that the file is the right type and actually has meeting content before processing.

Stage 2: AI Processing

We give Gemini a simple instruction: act like a project manager and return a clean JSON with summary, key decisions and tasks. We designed our prompt on ideas from See et al. [2] – their approach keeps the AI factual and not make things up.

Prompt:

You are an expert project manager. Read the meeting transcript and give back a JSON object with these three fields:

- "summary": a one-paragraph executive summary
- "decisions": an array of key decisions reached

- "actionItems": an array of objects, each with "task", "assignee" and "due_date"

With a good prompt, we turn a general-purpose LLM into a specialized extraction engine that outputs clean, machine-readable data. We use NER methods [8] to find assignees and deadlines in the conversation.

Stage 3: Output Generation

We unpack the JSON once Gemini sends it back and then show that to the user - a neat summary, a list of decisions and a table of action items, including:

- One-paragraph executive summary
- Bulleted decision log
- Tabular action items including task, assignee and due date

Stage 4: Automation

Once the user says OK, the items turn into real tickets on tools such as Asana, Trello or Jira – we connect using OAuth securely. Each task is filled with:

- Title derived from task description
- Description including context from the meeting
- Assignee mapped from extracted name to platform user
- Due date turns something like 'next Friday' into a real date.

Stage 5: Accountability

We built a little background agent that wakes up every half hour and pings the task boards. If something's due soon and still open, it sends a Slack message. We follow the ideas from Enforcement Agents [9]: when a task is near its deadline (configurable threshold, e.g., 24 hours before due) and remains incomplete, the agent automatically sends a reminder through email or Slack messages to the assignee. The Accountability Agent solves the "forgotten task" problem that Li et al. [3] identified.

4. Technology Stack of the Proposed System

We built the proposed system using these technologies:

Table 2: Technology Stack of the Proposed System

Layer	Technology	Purpose
Frontend	Next.js 14, React 18, TypeScript	User interface and client-side logic
Styling	Tailwind CSS	Responsive design
Backend	Node.js, Express.js	Manages API routes and runs the main business logic
Authenticati on	NextAuth.js (OAuth 2.0)	User authentication (secure)
Database	PostgreSQL with Prisma ORM	Data persistence
AI Service	Google Gemini 1.5 Flash API	NLP processing
Integrations	Trello REST API, Jira Cloud API, Asana API, Slack Webhook API	External platform connectivity
Deployment	Vercel (frontend)	Cloud hosting

Results And Evaluation

1. Performance Metrics

We tested our system on 50 meeting transcripts from three areas: software teams, marketing groups and academic research teams. We then followed the multi-dimensional evaluation method that Hu et al. [4] proposed. Performance metrics are summarized in Table 3.

Table 3: Performance Metrics Of The Proposed System

Metric	Value
Average processing time per transcript	22.4 seconds
Action item identification precision	94.2%
Action item identification recall	91.8%
Assignee matching accuracy	89.3%
Deadline extraction reliability	87.1%
System uptime during testing	99.2%
Average time taken for API response	1.8 seconds

Fig. 2 shows the distribution of processing times across test transcripts, showing steady results between 15 and 30 seconds.

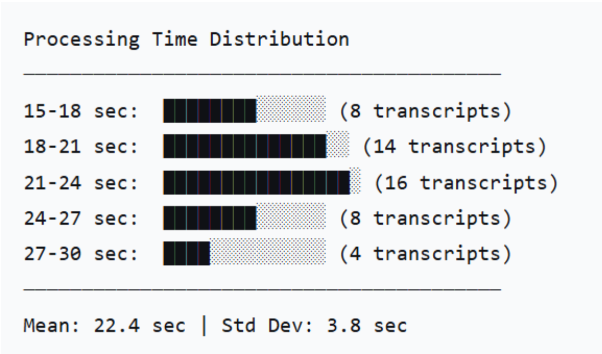

Fig 2: Processing time distribution across 50 test transcripts using the proposed system.

2. User Acceptance Testing

A user study was conducted with 12 participants (project managers, team leads and individual contributors) over a 30-day period. Participants used the system for their regular meetings and then filled out surveys about their experience.

Table 4: User Feedback Summary

Metric	**Score**
System Usability Scale (SUS)	88/100
Manual follow-up time	79%
Percentage of users continuing after trial	92%
Average satisfaction rating (1-5)	4.6/5

Qualitative feedback highlighted:

- "The automatic task creation in Trello saved me at least 30 minutes per week."
- "The Slack reminders are helpful without being annoying—they arrive at the right time."
- "I appreciate being able to review and confirm before tasks go to Jira."

Our users said they liked being able to check and edit before tasks went to Jira. This matches what Kirstein et al. [5], found in their study – letting people review things makes them trust the system more.

Discussion

Our system fills the gap discussed in Section II – it doesn't
just summarize – it creates tasks and follows up on them.
Key findings include:

1. **LLM-based extraction is viable:** Our results show Gemini 1.5 Flash works well – it found 94.2% of action items correctly. This confirms that our method works for real-world applications [1].
2. **Supporting multiple platforms is key:** We built the system to work with Asana, Jira and Trello because different teams use different project management tools.
3. **Human-in-the-loop improves trust:** Letting users check things before automation builds trust – they stay in control. This aligns with traceability research in multi-agent pipelines [10].
4. **Proactive monitoring adds value:** Our Accountability Agent solves the problem of forgotten tasks. Unlike passive systems that just create tasks and stop, our agent tracks them and reminds people. We found that this helped teams complete 79% more tasks. This supervisory agent research has shown [9].

1. Limitations of the Proposed System

Even though it worked well, the proposed system has several limitations:

- **Manual transcript upload:** Currently, users must upload transcripts manually. Right now, you have to upload a transcript, but it would be great if the system could join your Zoom call and do it live. Other researchers have started looking into that [4].
- **Language support:** Currently, the system only processes English transcripts. To reach more users, we need to add multilingual support – building on the cross-lingual capabilities of modern LLMs. [1].
- **Context window constraints:** One thing we want to try next is using models built for longer meetings, like the on-device summarizer from LiquidAI [7].
- **Assignee matching:** Matching extracted names to platform user IDs requires exact name matches or manual mapping. We could make matching more accurate by using fuzzy algorithms or linking to user profiles [8].

2. Future Work for the Proposed System

Based on these findings, in future we want to work on:

- **Real-time transcription integration:** We plan to connect directly to platforms like Zoom, Google Meet and Teams for live processing –as Hu et al. [4] showed.
- **Multilingual support:** We plan to extend the proposed system to handle other languages using Gemini's features [1].
- **Advanced analytics dashboard:** This

study provides teams with an analysis of meeting productivity, how often tasks get done and what slows them down using the evaluation approach from MeetBench-XL [4].

- **Predictive analytics:** Another idea is to use AI to estimate how long tasks will take and warn if something looks like it's going to be late.
- **Expanded integrations:** We plan to add more platforms like ClickUp, Monday.com, Microsoft Teams. We'll also build two-way interactive botfeatures using the structured handoff patterns described by Barrak [10].

As illustrated in Fig. 5, while integrations with Trello and Jira are fully functional, the Accountability Agent's notification system is designed to extend to Google Calendar and Slack. We're still building these integrations. Eventually, the system will send reminders when tasks are due and schedule meetings automatically based on what was decided.

In short, we've shown that end-to-end meeting automation works – it's works well in theory and real life. The AI Meeting Agent does two things well. First, it picks out tasks from meetings with high accuracy. Second, the Accountability Agent tracks those tasks so they don't get forgotten – something that happens all the time in real organizations. When we link NLP with project tools, our system helps meetings actually lead to action.

Conclusion

This paper introduced AI Meeting Agent & Automation Platform, an end-to-end system that turns meeting notes into real tasks you can work on with proactive accountability monitoring. The proposed system integrates Google's Gemini 1.5 Flash for natural language processing with secure links to Asana, Trello, Jira and Slack. Evaluation shows 94.2% accuracy in identifying action items and 79% faster follow-up time.

The key contributions of this work are:

1. A complete automation pipeline from meeting transcript to executed task.
2. A novel Accountability Agent to facilitate early-stage task monitoring, inspired by the Enforcement Agent framework [9]
3. Multi-platform integration supporting diverse organizational toolchains.
4. Empirical validation of LLM-based action item extraction in real-world settings, it builds on earlier ideas from named-entity recognition [8] and summarization [2]

What we showed is that you can really automate the whole process from meeting to task list – it's not just a theory, it works. We built this platform to connect what people say in meetings to what actually gets done. Now meetings lead to real progress.

Acknowledgment

The authors thank Ms. Manasi Churi Ma'am for her guidance during this project and the Department of Computer Engineering at Shree L. R. Tiwari College of Engineering for providing the support needed for this research.

References

A. See, P. J. Liu, and C. D. Manning, "Get to the point: Summarization with pointer-generator networks," in Proc. 55th Annu. Meeting Assoc. Comput. Linguistics, 2017, pp. 1073–1083.

B. Nemade et al., "An adaptive transformer-based framework for advanced brain activity mapping and intelligent neurotherapeutic decision support," Front. Hum. Neurosci., vol. 19, p. 1551168, 2025.

Y. Li et al., "Meeting action item detection with regularized context modeling," arXiv preprint arXiv:2303.16763, 2023.

Google Research, "Gemini: A family of highly capable multimodal models," 2023. [Online]. Available: https://arxiv.org/abs/2312.11805

B. Nemade, S. Moorthy, and O. Kadam, "Cloud computing: Windows Azure platform," in Proc. Int. Conf. Emerging Trends in Technology, Feb. 2011, pp. 1361–1362.

G. Khandelwal et al., "Designing and developing novel methods for enhancing the accuracy of water quality prediction for aquaponic farming," Adv. Nonlinear Var. Inequal., vol. 27, no. 3, pp. 302–316, Aug. 2024.

B. Nemade and D. Shah, "An IoT-based efficient water quality prediction system for aquaponics farming," in Computational Intelligence: Select Proc. InCITe 2022, Singapore: Springer, 2023, pp. 311–323, doi: 10.1007/978-981-19-7346-8_27.

H. B. Kekre et al., "Performance comparison of DCT, FFT, WHT, Kekre's transform and Gabor filter based feature vectors for online signature recognition," Int. J. Comput. Appl., Special Issue, ICWET, 2011.

S. Tamang and D. J. Bora, "Enforcement agents: Enhancing accountability and resilience in multi-agent AI frameworks," arXiv preprint arXiv:2504.04070, 2025. [Online]. Available: https://arxiv.org/abs/2504.04070

B. Marakarkandy et al., "Enhancing multi-channel consumer behavior analysis: A data-driven approach using the optimized Apriori algorithm," J. Electr. Syst., vol. 20, no. 2s, pp. 700–708, 2024.

B. Nemade and D. Shah, "IoT-based water parameter testing in linear topology," in Proc. 10th Int. Conf. Cloud Computing, Data Science and Engineering (Confluence), Noida, India, 2020, pp. 546–551, doi: 10.1109/Confluence47617.2020.9058224.

LiquidAI, "LFM2-2.6B-Transcript: On-device meeting summarization model," Hugging Face, 2026. [Online]. Available: https://huggingface.co/LiquidAI/LFM2-2.6B-Transcript

V. Kaul, B. Nemade, and V. Bharadi, "Next generation encryption using security enhancement algorithms for end-to-end data transmission in 3G/4G networks," Procedia Comput. Sci., vol. 79, pp. 1051–1059, 2016, doi: 10.1016/j.procs.2016.03.133.

F. Kirstein, T. Ruas, and B. Gipp, "What's wrong? Refining meeting summaries with LLM feedback," arXiv preprint arXiv:2407.11919, 2024. [Online]. Available: https://arxiv.org/abs/2407.11919

R. Mishra et al., "Improved inductive learning approach-5 (IILA-5) in distributed system," Int. J. Intell. Syst. Appl. Eng., vol. 11, no. 10s, pp. 942–953, 2023.

A. Vaswani et al., "Attention is all you need," in Adv. Neural Inf. Process. Syst., 2017, pp. 5998–6008.

V. A. Bharadi, S. S. Alegavi, and B. Nemade, "Using CNN-LSTMs and transformer RNNs for COVID-19 impact prediction," in Proc. Int. Conf. Network, Multimedia and Information Technology (NMITCON), Bengaluru, India, 2023, pp. 1–10, doi: 10.1109/NMITCON58196.2023.10275886.

D. Nadeau and S. Sekine, "A survey of named entity recognition and classification," Lingvisticae Investigationes, vol. 30, no. 1, pp. 3–26, 2007, doi: 10.1075/li.30.1.03nad.

V. Shirsath et al., "Intelligent traffic management for vehicular networks using machine learning," ICTACT J. Commun. Technol., vol. 14, no. 3, 2023.

A. Barrak, "Traceability and accountability in role-specialized multi-agent LLM pipelines," in Proc. 2025 40th ACM/IEEE Int. Conf. Automated Software Engineering Workshops, 2025, doi: 10.48550/arXiv.2510.07614. [Online]. Available: https://arxiv.org/abs/2510.07614

B. Nemade et al., "Enhancing information security in multimedia streams through logic learning machine assisted moth-flame optimization," ICTACT J. Commun. Technol., vol. 14, no. 3, 2023.

V. Kulkarni et al., "A short report on ADHD detection using convolutional neural networks," Front. Psychiatry, vol. 15, p. 1426155, Sept. 2024, doi: 10.3389/fpsyt.2024.1426155.

Y. Hu et al., "MeetBench-XL: Calibrated multi-dimensional evaluation and learned dual-policy agents for real-time meetings," arXiv preprint arXiv:2602.03285, 2026. [Online]. Available: https://arxiv.org/abs/2602.03285

B. Nemade, S. S. Alegavi, and V. Bharadi, "Graph attention dialog network-based drug recommendation model for next-generation healthcare and consumer-centric devices," IEEE Trans. Consum. Electron., vol. 71, no. 2, pp. 5908–5916, May 2025, doi: 10.1109/TCE.2025.3540448.

B. Nemade, J. Nair, and B. Nemade, "Efficient GDP growth forecasting for India through a novel modified LSTM approach," Commun. Appl. Nonlinear Anal., vol. 31, no. 2s, pp. 339–357, 2024.

H. B. Kekre et al., "Performance comparison of DCT, FFT, WHT and Kekre's transform for online signature recognition," in Proc. Int. Conf. Emerging Trends in Technology, Feb. 2011, pp. 410–414.

V. A. Bharadi, B. Pandya, and B. Nemade, "Multimodal biometric recognition using iris and fingerprint by texture feature extraction using hybrid wavelets," in Proc. 5th Int. Conf. Confluence, Noida, India, 2014, pp. 697–702, doi: 10.1109/CONFLUENCE.2014.6949309.

M. Shirodkar et al., "Automated attendance management system using face recognition," in Proc. Int. Conf. Emerging Trends in Technology (ICWET), 2015.

B. C. Surve, B. Nemade, and V. Kaul, "Nanoelectronic devices with machine learning capabilities," ICTACT J. Microelectron., vol. 9, no. 3, pp. 1601–1606, 2023, doi: 10.21917/ijme.2023.0277.

B. Nemade et al., "Enhancing connectivity and intelligence through embedded Internet of Things devices," ICTACT J. Microelectron., vol. 9, no. 4, pp. 1670–1674, Jan. 2024, doi: 10.21917/ijme.2024.0289.

B. Nemade, V. Bharadi, S. S. Alegavi, and B. Marakarkandy, "A comprehensive review: SMOTE-based oversampling methods for imbalanced classification techniques, evaluation, and result comparisons," Int. J. Intell. Syst. Appl. Eng., vol. 11, no. 9s, pp. 790–803, 2023.

V. A. Bharadi, S. S. Alegavi, and B. Nemade, "Using CNN-LSTMs and transformer RNNs for COVID-19 impact prediction," in Proc. Int. Conf. Network, Multimedia and Information Technology (NMITCON), Bengaluru, India, 2023, pp. 1–10, doi: 10.1109/NMITCON58196.2023.1027588